

THE CHRONICLE OF HIGHER EDUCATION

Is This Economist Too Far Ahead of His Time?

David Wescott

October 16, 2016

It's the year 2120. You feel no hunger, no cold, no heat, no pain. There's no need to eat or to take medicine, though you can if you like. You are beautiful, intelligent, and charismatic, as are your friends, co-workers, lovers. Though the economy is fiercely competitive, retirement is not far off. You do not fear death. Look out your office window and you see sunlit spires towering over tree-lined boulevards.

At least this is what you think you see. In fact, you live and work in virtual reality. Your city amounts to racks of computer hardware and the pipes that cool them. And you are not "you" in the traditional sense: You are an "em," a robotic brain emulation created by scanning a particular human brain and uploading it to a computer. On the upside, you process information 1,000 times faster than a human. On the downside, you inhabit a robotic body, and you stand roughly two millimeters tall.

This is the world Robin Hanson is sketching out to a room of baffled undergraduates at George Mason University on a bright April morning. To illustrate his point, he projects an image of an enormous futuristic city alongside clip art of a human castaway cowering on a tiny desert island. His message is clear: The future belongs to "ems."

This may sound more like science fiction than scholarship, but that's part of the point. Hanson is an economist with a background in physics and engineering; a Silicon Valley veteran determined to promote his theories in an academy he finds deeply flawed; a doggedly rational thinker prone to intentionally provocative ideas that test the limits of what typically passes as scholarship. Those ideas have been mocked, mimed, and marveled at — often all at once.

Hanson, deeply skeptical of conventional intellectual discourse, argues that academics have abdicated their societal responsibilities by ignoring more speculative work. "Relative to the future, our study of the past has hit diminishing returns," he writes in his first book, *The Age of Em: Work, Love, and Life When Robots Rule the Earth* (Oxford University Press), published this year. It challenges readers' expectations for scholarly work, arguing that an insistence on considering only those ideas with strong supporting evidence needlessly discards useful thinking.

While his ideas have raised plenty of eyebrows, Hanson is not quite an academic pariah. Wendell Wallach, a scholar at Yale University's Interdisciplinary Center for Bioethics, describes him as "truly a unique thinker" for the way he uses social science and economics to try to prognosticate the future. "That's very different than what everybody else does," he says. Nick Bostrom, a professor of philosophy at Oxford and director of the Future of Humanity Institute (with which Hanson has an affiliation), calls him "perhaps the most original thinker out there in the social sciences today."

After the lecture, Hanson decamps to Carrow Hall, a one-story building on George Mason's suburban Fairfax, Va., campus, which houses the Center for Study of Public Choice. In a meeting area near his book-crammed office, he launches into a characteristically sweeping riff on the nature of the university.

"One of the signature problems about academia is that it seems to clump," he says. "There's this vast space of potentially interesting topics, and you see academic clump on the same few topics." Research done in a certain subfield legitimizes the subfield and paves the way for similar research, he reasons. "The space between the clumps gets neglected."

Hanson, 6 feet tall with thinning light-brown hair, talks rapidly and often delves into dialogue with himself. He's fond of thought experiments that proceed logically but end up somewhere unexpected, and seems distrustful of his own intuition. "A key fact about being a weird person is puzzling at how the world seems so different from you," he explains. "I am a weird person, and I puzzle. That means I have to be very reluctant to use my own intuitions about things as a guide."

Like much of the George Mason economics department, Hanson leans libertarian, but he has also dreamed up his own form of government, called "futarchy." An enthusiastic student describes his true political philosophy as "meta, meta-everything": questioning how we arrive at political philosophies in the first place. To the dismay of his wife, Peggy Jackson, who has worked in hospice care, he is a member of the Alcor Life Extension Foundation, paying a monthly fee so that one day his head will be cryogenically frozen. He thinks he may live a long, long time — if not forever, then possibly at least long enough to participate in the em future.

Hanson takes ideas, and their consequences, seriously, and he thinks a lot of scholars don't. He abhors self-deception and hypocrisy — which, in his view, arise when people aren't extremely open, honest, and forthcoming. Given this definition, he sees hypocrisy nearly everywhere. A persistent problem with academe, to his mind, is its focus on signaling impressiveness for its own sake. He is critical of the type of articles that top journals tend to look for: "It will have to use difficult methods, more-difficult-to-access data sets, more-difficult-to-access theoretical concepts," he explains. The major criterion is impressiveness, or difficulty, "not actual intellectual contribution or insight." In 1990 he wrote that "academia is still largely a medieval guild, with a few powerful elites, many slave like apprentices, and members who hold a monopoly on the research patronage of princes and the teaching of their sons."

Hanson considers himself something of an exception to that rule and has described his mission as a "sacred quest, to understand everything, and to save the world." He argues that academics are primarily devoted to signaling their own importance, and not necessarily to the pursuit of intellectual progress. "We lie about why we go to prestigious colleges as students, we lie about why we fund research, we lie about why we do research ... we lie about lots of things," he says. "We are so tempted to bullshit and give the most noble reason for why we do things."

For academics who do actually care about intellectual progress more than "prestige, promotions, salaries, funding, lots of students, and roaring crowds," Hanson says, there is a lot of freedom.

For him it's the freedom to study things like immortality, aliens, and what to do if you suspect you are living in a computer simulation. "There are important silly subjects," he says. And while most academics shy away from silly, "silly doesn't equal unimportant."

He knows he's lucky to have "skirted by the edge" in his career. "I'm a successful enough academic to have tenure while violating many pieces of good advice to young people about what not to do."

Robin Hanson was born in Chicago in 1959 to devout Baptist parents. On weekends they would drive to the countryside to hold evangelical services; Hanson and his two brothers were enlisted as the choir, decked out in matching outfits. When he was 12 or 13, he joined a religious cult. The cult believed that the end times were imminent. His parents found out about the group after a while and put an end to his participation. To this day, Hanson has a fascination with utopian communal living groups.

In high school in California, Hanson discovered physics, which offered a comprehensive explanation of the world and displaced religion in his life. He enrolled in the engineering school at the University of California at Irvine and bounced between engineering and physics, finding neither particularly challenging. He stopped doing his homework and just played around with equations. He aced his exams anyway.

In 1981 he went to the University of Chicago to pursue master's degrees in physics and the philosophy of science. He did well academically but grew increasingly impatient. Hanson came to believe that some big questions in physics, like those in quantum mechanics, already had answers, like the many-worlds theory, which suggests that all alternative universes are real — and which Hanson estimates to be 95 percent likely. He had little interest in refining established theories and preferred unexplored areas where he could focus on big, neglected questions.

Two papers he read during this period stayed with him: one on hypertext publishing and one on artificial intelligence. To pursue these ideas, he moved to Silicon Valley and got a job at

Lockheed Martin in statistics and AI research. In his free time, he worked with Project Xanadu, an experimental effort to build an advanced hypertext system.

Project Xanadu was in part an attempt to improve scholarly debate. The idea was for hypertext links to present contrary evidence or useful information. The best evidence would be impossible to ignore, and the world of ideas would be increasingly governed by rational exchange, supported by data.

At the same time, Hanson himself felt isolated from the world of ideas. Without a Ph.D., he felt he wasn't being taken seriously. So, at 34, with two young children and a wife who wanted to stay home with them, he quit his job and enrolled in the doctoral program in social science at the California Institute of Technology. A "heroic Hail Mary pass on my career," he calls it.

At Caltech he found that economists based their ideas on simple models, which worked well in experiments but often failed to capture the complexities of the real world.

That became painfully clear to him about a decade later, in 2003, when Hanson, then at George Mason, and John Ledyard, an economist who had been his supervisor at Caltech, joined with several others to work on a proposal for Darpa, the Defense Advanced Research Projects Agency. It was a project to build a betting market for ideas, or "idea futures." They called it the Policy Analysis Market.

Hanson had proposed the concept as early as 1988. With money on the line, he believed, scholars or policy analysts would no longer be able to bluster, seek merely to sound smart, or embrace contrarianism for contrarianism's sake. They would be motivated to champion the best ideas, without cynicism or hypocrisy.

The project was going well — experts were within months of being able to place bets of less than \$100 on the likelihood of things like civil unrest in Saudi Arabia. Then Democratic senators objected and the news media caught wind of the project, excoriating it as a "terrorism futures market" that might allow terrorists to profit from carrying out their missions.

The Policy Analysis Market fell apart, and Ledyard and Hanson were left to consider what might have been. In typical fashion, Hanson conducted a statistical model of the news coverage. His finding: "The Informed Press Favored the Policy Analysis Market."

The failure wasn't a career killer. In fact, Hanson wonders if being denounced by Democratic senators helped him win tenure at George Mason.

In *The Age of Em*, Hanson takes what he calls the "somewhat unusual approach of using basic social theory, in addition to common sense and trend projection, to forecast future societies." He allows that the book analyzes doomsday scenarios that others would eschew: "It assumes that our social systems will mostly fail to prevent outcomes that many find lamentable, such as robots dominating the world."

The book's premise is that in about a century, it will be possible to scan human brains at "fine enough spatial and chemical resolution," and to combine that with good-enough models of signal-processing functions of brain cells, "to create a cell-by-cell dynamically executable model of the full brain in artificial hardware, a model whose signal input-output behavior is usefully close to that of the original brain."

Hanson draws substantially from the research of others, particularly the "Whole Brain Emulation" Roadmap, a 2008 paper by Anders Sandberg and Nick Bostrom, of Oxford's Future of Humanity Institute. While Hanson admits that there's less than a one-in-a-thousand chance of his predictions' coming entirely true, he estimates, conditionally, that 30 percent of possible future realities will be usefully informed by his analysis.

Some scholars are sympathetic to this view. As Wallach, the Yale technology ethicist, puts it, "I think of the next 50 to 100 years as science fiction. But even science fiction informs us."

In that context, a scenario that may be useful for 30 percent of likely future realities is one worth exploring. And explore Hanson does, with a sheer force of detail that reads like a bizarre history textbook.

He suggests that most ems will be copies of the same 1,000 or so "best" human brains; everyone else will end up like the cowering castaway he showed his students. The ems themselves will be easily copied, which will lead to ems creating "spurs" of themselves — short-lived versions that may do an unpleasant task and then come to an end. Instead of worrying about death, ems will ask themselves: "Do I want to remember this?" If not, turning oneself off is no great loss.

Em society will be organized in clans, with rival factions competing and jealously guarding their intellectual property — their minds. While death isn't feared, "mind theft" will be.

The em economy will grow so quickly — Hanson estimates it will double every month or so — that this chapter in the society's existence will be short-lived. Perhaps this age of ems will last a year or two, although it will feel much longer, subjectively, to the high-speed-processing ems. It will then give way to something stranger, something outside even Hanson's purview. But what kind of economist thinks all of this is within his purview in the first place?

This idea of robotic brain emulations is not unique to Hanson — other scholars study em scenarios, dwelling on questions like, "Would ems be conscious?" or "Would it really be me?" For Hanson, those philosophical ponderings are beside the point. Whether we like it or not, and whether ems would be "alive" or not, he believes, this future could very well come to pass.

A recent model for *The Age of Em* is Jørgen Randers's *2052: A Global Forecast for the Next Forty Years* (Chelsea Green, 2012). Randers is a professor emeritus of climate strategy at the BI Norwegian Business School and an author of the influential 1972 book *The Limits to Growth*. "In academia, even today, science is organized in silos," he says. "In order to get ahead in the math silo, you have to specialize like hell in a mathematical discipline. This goes in political science, sociology, economics, even in biology." By contrast, Randers, like Hanson, endorses a systems perspective. "It's not broad and deep, because no one can be deep all across the whole spectrum. Seen from the vantage point of the specialist, any type of systems analysis or forecast of the world appears to be shallow."

And it's hard to get funding for projects deemed shallow. "Boards have to send these applications for peer review. They send, of course, to the historian, the biologist, the mathematician. Each of them comes back and says, 'This is an interesting study, but in my niche, it ought to do such and such.'"

"Why can you do scholarship on 100 years ago and not 100 years from now?" asks Max Tegmark, a professor of physics at the Massachusetts Institute of Technology and president of the Future of Life Institute. "Within each discipline it should be as legitimate to study the future as it is to study the past."

Intradisciplinary reform may not be in the offing, but an influx of money from Silicon Valley is mitigating the challenges to futurist research in the academy. The Future of Humanity Institute, at Oxford, was started with a large donation from James Martin, a technology writer and Oxford alumnus. "Once you have your own island in Bermuda, what do you do? You try to save the world," says Sandberg, a research fellow at the institute. Elon Musk has given \$10 million to Tegmark's Future of Life Institute, which funds research on AI.

Most of science works "by taking ridiculously small steps where they check everything a hundred times over and slowly progress," says Sandberg, who labels most scientists "very myopic." They simply don't ask, "Where are we heading, in the large?"

"Let's not drive down the highway with our eyes closed," says Tegmark. The blind pursuit of technology, without pursuing the equivalent wisdom in how to guide that technology, is, he thinks, a recipe for disaster.

Evaluation is another challenge for academic futurists. How can we assess a proposed future of Lilliputian brain emulations? "It took a century to verify some of Einstein's theories," says Erik Brynjolfsson, a professor at MIT's Sloan School of Management. While "it's great when scientific theories can be evaluated against empirical evidence," he says, "the reality is that there are a lot of theories that it's hard to have clear empirical evidence for one way or the other." String theory, for example.

Brynjolfsson thinks we need more economists like Hanson, and he encourages his students not to do what everyone else is doing, but to pursue riskier work. "You need a spectrum of people," he says, invoking the views on entrepreneurship of Peter Thiel, the venture capitalist. "We could use more people who go against the grain, outside of the mainstream, who don't follow the crowd."

Not following the crowd has led Hanson to largely empty corners of the academy, but no one said the "sacred quest, to understand everything, and to save the world" would be easy — or lucrative. He received a \$1,000 advance for *The Age of Em* and points out that he's the lowest-paid tenure-track economics professor at George Mason. (Both of his twenty-something kids earn more than he does.)

If his out-there intellectual pursuits have not lead to great riches, it is not due to a lack of zeal in promoting his vision. In addition to sending a draft of the book around to 500 intellectuals, academics, and futurists for comment before publication, he's printed 1,000 index card-sized advertisements to hand out at the more than 75 talks he's giving to promote *The Age of Em*.

One such event takes place on a Tuesday morning in early June at the Cato Institute, a libertarian think tank in Washington. Wearing a pink shirt and red tie, with cream-colored pants and jacket, he's a bit out of sorts — traffic was terrible — but soon hits his stride. About 45 people attend, including professors, interns from a range of libertarian organizations, and miscellaneous opportunists there for the free lunch that follows. There's talk of Matrioshka brains, transcension, and living in blocks of computronium.

Hanson sells only one book. It doesn't help that Cato can't take credit cards.

After the talk, Hanson deposits a box of unsold books in his car (a bumper sticker reads "Question Authority: But Raise Your Hand First"), and we make our way to a coffee shop. We talk about the merits of writing a book versus those of blogging. Hanson is a prolific blogger at Overcoming Bias, which receives tens of thousands of page views per month. There, over the course of thousands of posts (more than 250 of which are labeled "Academia") he's made the case for his style of thinking. Among his provocations there: He has argued that we should forget 9/11, and, via thought experiment, wondered why society punishes instances of "gentle silent rape" more severely than it does those of cuckoldry. (He later added a trigger warning to the post but didn't take back the incendiary phrasing.)

The conversation turns to the futurist's own future. He's co-writing a book on bias and hypocrisy for Oxford University Press. He has a three-year grant from the Open Philanthropy Project to study conventional AI scenarios (in which software-driven machine learning reaches human-level intelligence before brain emulations get there). He also thinks about writing fiction based on *The Age of Em* — he has a drawerful of ideas.

A few times a week, Hanson has lunch with a group of George Mason economists. One brisk Thursday, Tyler Cowen, John Nye, Bryan Caplan, Alex Tabarrok, and Hanson drive over to China Star, a Sichuan-style Chinese restaurant near the campus.

The atmosphere is collegial; they talk about bets they're making with one another, what they've already won and lost (also painstakingly detailed online, as several of them are prominent bloggers), and tease Hanson for his grandiose visions of immortality. Cowen, author of *An Economist Gets Lunch* (Dutton, 2012), orders for the table, and over spicy dishes passed around they dig in.

"Robin is so fond of generalizations that he'll often ignore varying details," Nye says. "It leads to good ideas, but it also leads to, in my view, crazy ones."

"That comes from physics," replies Hanson.

"There's a reductionism that comes from physics," Cowen says. "Reductionism, monism, trying to recreate the problems of theology, moralizing, 'meta,' and hating hypocrisy — that's Robin in 10 words. He's a modern gnostic."

Later, after digressions into *The Lord of the Rings*, futarchy, and the relative innovativeness of the iPhone, Hanson wonders aloud why his ideas aren't more widely circulated or accepted in academe. His colleagues don't hesitate to offer theories.

"Robin's work would be much more accepted if he just did one weird thing and everything else was normal," says Caplan. "If everything was normal but he did the future, that'd be OK. But he has seven or more weird things."

"I'm rolling more dice, so there's more of a chance one of them will come out right," Hanson says.

"But," replies Nye, "there's also more of a chance it's crazy."